Certified and fast computation of supremum norms of approximation errors

Sylvain Chevillard, Mioara Joldes, Christoph Lauter

Labo de l'Informatique du Parallelisme, 46 Allee d'Italie, Lyon, France

In many numerical programs there is a need for a high-quality floating-point approximation of useful functions f, such as such as exp, sin, erf. In the actual implementation, the function is replaced by a polynomial p, which leads to an approximation error (absolute or relative) $\epsilon = p - f$ or $\epsilon = \frac{p}{f} - 1$. The tight yet certain bounding of this error is an important step towards safe implementations.

The problem is difficult mainly because that approximation error is very small and the difference $p - f$ is subject to high cancellation. Previous approaches for computing the supremum norm in this degenerate case, have proven to be unsafe, not sufficiently tight or too tedious in manual work.

We present a safe and fast algorithm that computes a tight lower and upper bound for the supremum norms of approximation errors. The algorithm is based on a combination of several techniques, including enhanced interval arithmetic, automatic differentiation and isolation of the roots of a polynomial.

Moreover, we emphasize the problem of high dependency present in the case of approximation errors (relative and absolute) and how this problem can be alleviated using techniques based on automatic differentiation or Taylor models. We compare these methods for the problematic case of relative errors.

Finally, we give the timings obtained with our algorithm on several examples.